

MECHANISM FOR DATA AGGREGATION IN A TRACING FRAMEWORK

Abstract

A method is disclosed for obtaining data from a kernel, including obtaining data from the kernel, and storing the data in a data set an aggregation buffer using an aggregation function. A method is disclosed for storing data in a data set, wherein the data set includes a key component, an aggregation identifier component, and a value component, including obtaining an expression, a new value, and an aggregation identifier, generating a key using the expression and the aggregation identifier; and storing the data set in a buffer, wherein storing the data set comprises storing the key in the key component, storing the aggregation identifier in the aggregation identifier component, and updating a current value in the value component using the new value and an aggregation function.

55918_3.DOC